STAFF SUMMARY SHEET

	TO	ACTION	SIGNATURE (Surname), GRADE AND DATE		TO	ACTION	SIGNATURE (Surname), GRADE AND DATE
1	DFEG	sig	MACHOVINA O-5, 6 OCT 14	6			
2	DFER	approve	Sofft, AD22, 6 Oct 14	7			
3	DFEG	action	Maj Ackerman	8			
4				9			
5				10			

SURNAME OF ACTION OFFICER AND GRADE	SYMBOL	PHONE	TYPIST'S INITIALS	SUSPENSE DATE
Maj Ackerman	DFEG	9283	ada	20141009

SUBJECT: Clearance for Material for Public Release — USAFA-DF-PA- 455

DATE: 20140827

SUMMARY

1. PURPOSE. To provide security and policy review on the document at Tab 1 prior to release to the public.

2. BACKGROUND.
Authors: Adam Ackerman, Elizabeth Carpenter, Sierra Kelly, Robert Gutierrez

Title: Effectiveness of Incorporating Adversary Probability Perception Modeling in Security Games

Circle one: Abstract Tech Report Journal Article Speech (Paper) Presentation Poster

Thesis/Dissertation Book Other: _____

Check all that apply (For Communications Purposes):

[] CRADA (Cooperative Research and Development Agreement) exists

[] Photo/ Video Opportunities [] STEM-outreach Related [] New Invention/ Discovery/ Patent

Description: This paper will be submitted to the 2015 Association for the Advancement of Articial Intelligence (AAAI) Symposium on Applied Computational Game Theory. Submissions are due 10 Oct 2014.

Release Information: See document header. needs disclaimer

Previous Clearance information:

Recommended Distribution Statement: Distribution A: approved for public release, distribution unlimited

3. DISCUSSION.

4. RECOMMENDATION.

ADAM D. ACKERMAN, Major, USAF
Instructor of Economics

Tab(s)
1. Effectiveness of Incorporating Adversary Probability Perception Modeling in Security Games

This research was performed under an appointment to the US Department of Homeland Security (DHS) Science & Technology (S&T Directorate Office of University Programs Summer Research Team Program for Federal Service Academies, administered by the Oak Ridge Institute for Science and Education (ORISE) through an interagency agreement between the US Deparment of Energy and DHS. ORISE is managed by ORAU under DOE contract number DE-AC05-06OR23100. All opinions expressed in this paper are the authors' and do not reflect the policies and views of DHS, DOE, ORAU/ORISE, United States Air Force, Department of Defense, or the United States Government.

Effectiveness of Incorporating Adversary Probability Perception Modeling in Security Games

Adam Ackerman, Elizabeth Carpenter, Sierra Kelly, Robert Gutierrez
Department of Economics and Geosciences
United States Air Force Academy, CO 80840
adam.ackerman@usafa.edu

Abstract

We attempt to further improve the Subjective Utility Quantal Response (SUQR) model used to handle boundedly rational human adversaries in Stackelberg security game (SSG) algorithms. Given recent work on human decision-making, we adjust the existing subjective utility function to account for adversary probability perceptions of defender coverage. We then compare the predictive accuracy of this expanded model to the existing SUQR model utilizing data from previous security game experiments with human subjects. Our results show the incorporation of probability perceptions into the SUQR can provide improvements in the ability to predict probabilities of attack in certain games.

Introduction

The problem we address involves a SSG where a defender has limited security resources to deploy in order to protect valuable targets, and an attacker has incentives to hit these targets. The defender can cover each of the targets with some probability, and the attacker can choose to strike the targets which will provide the most benefit to them. The real-life applications of the Stackelberg security environment include gate security at airports, illegal fishing activity in coastal waters, and illegal poaching in protected forests. Security games have been successfully applied to aid security agencies to allocate their limited resources against potential threats (Pita et al 2008). The motivation behind this research is to improve upon existing models which predict attacker responses to defender strategies so that defenders can better deter attacks. By increasing costs for attackers, hopefully it will lead to less terrorist, poaching, and illegal fishing activity. To evaluate the model, we analyzed data from human subject experiments playing security games. We transformed the original model to include perceived probabilities. Then, we tested this expanded model against a previous model which did not include perceived probability weighting.

Our research focuses on improving the subjective utility function in the SUQR model which predicts probabilities of attack. The existing SUQR model is an advanced model that departs from the assumption that attackers are perfectly rational. Specifically, the model assumes "bounded rationality" which allows better prediction of more uncertain human behavior consistent with real world observations. Still, the previous work does not address an adversary's perception of defender coverage probability, in other words, the probability that an attacker thinks he will be caught. In reality, adversaries may underestimate or overestimate the true probability that they will be caught. Our work seeks to account for these probability perceptions. Probability weighting functions have not been applied to the SUQR model, so the question as to whether it will improve with their addition is yet to be answered. Our work could lead to an improved model which more accurately predicts attack probability because of increased weighting on probability perception.

We relied upon previously collected data from SSGs originally deployed on Amazon Mechanical Turk (AMT) for "Analyzing the Effectiveness of Adversary Modeling in Security Games" (Nguyen et al. 2013). This game data provides us potential attackers' strategies where the attacker is not completely rational.

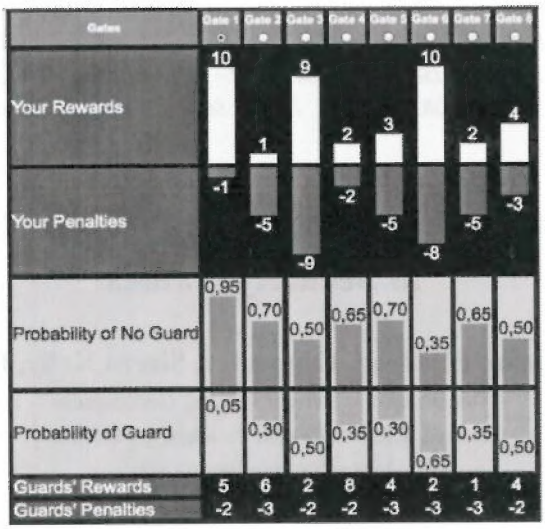

Figure 1. Game Interface

The games simulate the security environment at an airport with 8 or 24 gates (targets), which are protected by 3 or 9 guards, respectively. The Graphical User Interface (GUI) seen in Figure 1 displays the coverage probability (probability of a guard being present) along with the reward for a successful attack and the penalty for a failure at each gate. If the attacked gate is unguarded, the attacker receives a reward and the defender gets a penalty, and vice versa. For both the 8 and 24 gate games, we examined 44 separate payoff structures randomly generated from the multivariate normal distribution, and the presented defender coverage strategy was previously learned by the Subjective Utility Quantal Response (SUQR) model in the same setting, i.e. 8 or 24 gate and 3 or 9 guard games with AMT workers. Each participant had one chance to choose their point of attack in each game. An average of 45 participants, all from the the United States, played each of these games. For the 8 gate games, we utilized a total of 2086 observed attacks (1636 for training and 450 for testing). For the 24 gate games, we utilized a total of 1903 observed attacks (1508 for training and 395 for testing). After learning 'simple' SUQR and new 'expanded' SUQR models with perceived probabilities, we evaluate the different models with the real-world data in order to provide credible recommendations. Ideally, this work will lead to the application of more effective defender strategies for security in actual airports and other real-world environments.

Background and Related Work

In SSGs, a great deal of previous work has contributed to modeling the decisions of attackers as quantal response problems. This idea of quantal response simply means attackers have a quantifiable number of ways to attack. For example, an attacker in the 8-gate security game has 8 possible strategies. A practical approach to the specification of these decisions stems from Luce's choice model, often referred to as the strict utility model (McFadden 1976). When independence of irrelevant alternatives criterion are met, Luce's model explains that the relative odds of a specific choice remain the same even when additional choices are introduced (Luce 1959). In this model, all responses or choices have non-zero probabilities, but it is unlikely an individual will make an extremely low valued choice. In our domain of SSGs, this model can represent the quantal response of an adversary, or the probability of attack for a gate t as seen in equation 1. The t' notation represents $t=1$ to the total number of gates.

$$q_t = \frac{e^{t'a}_t}{\Sigma_{t'} e^{t'a}_{t'}} \qquad \text{(equation 1)}$$

In behavioral game theory, the concept of quantal response equilibrium suggests instead of strictly maximizing utility, individuals make errors in games (Mckelvey and Palfrey 1995). Luce's model was expanded upon by introducing a rationality parameter, λ, which represents the amount of noise in an attacker's response. A λ of 0 implies completely irrational play and

results in a uniform attack strategy, while a λ of infinity implies perfect rationality and a strategy that optimizes an attacker's utility. This model is represented below in equation 2 (Nguyen et al. 2013).

$$q_t = \frac{e^{\lambda U_t^a}}{\Sigma_{t'} e^{\lambda U_{t'}^a}} \quad \text{(equation 2)}$$

More previous research in the SSG domain has also investigated applying different utility functions to this model. Work initially examined using an expected utility function as represented in equation 3. In this equation, the defender coverage probability is represented by x_t while the adversary's rewards and penalties are represented by R_t^a and P_t^a respectively.

$$U_t^a = x_t P_t^a + (1 - x_t) R_t^a \quad \text{(equation 3)}$$

To account for attacker utility, researchers also examined another human behavior theory, the nobel-prize winning Prospect Theory, which suggests individuals make decisions via maximizing their 'prospect' (Kahneman and Tvesky, 1979). The 'prospect' can be represented as a function of the perceived probability and value of an outcome. The researchers adjusted the probability weighting and value functions of the above expected utility equation according to the behavioral theory. Surprisingly, the results found the model incorporating Prospect Theory never outperformed the QR model with the simpler expected utility function (Yang et al 2011).

Since the QR model outperformed the model applying Prospect Theory, further work focused on different utility functions which led to the improved SUQR model. This model combined, in a linear function, the information presented to the attacker about each target choice (i.e. potential rewards, penalties and target coverage). Equation 4 shows the 3-parameter version of this function.

$$U_t^a = w_1 x_t + w_2 R_t^a + w_3 P_t^a$$

(equation 4)

A possible explanation for the success of this function in accurately modeling utility is that humans use simple heuristics in their decision making (Nguyen et al. 2013). Another possible explanation stems from the Lens model, which has been proposed to explain the human judgement process as coming from cues (Hammond 1955). Specifically in this case, some combination of a multiplicity of cues, like potential rewards, penalties and target coverage, may reflect an adversary's likelihood to attack a target. In addition to the 3-parameter function introduced above, researchers also investigated the use of a 5-parameter function with defender rewards and penalties. They found the 3-parameter function performed better than both the 5-parameter function and the basic expected utility function (Nguyen et al. 2013). Equation 5 displays the incorporation of the subjective utility function into the probability of attack function for the model (Nguyen et al. 2013).

$$q_t = \frac{e^{\lambda U_t^a}}{\Sigma_{t'} e^{\lambda U_{t'}^a}} = \frac{e^{\lambda(w_1 x_t + w_2 R_t^a + w_3 P_t^a)}}{\Sigma_{t'} e^{\lambda(w_1 x_{t'} + w_2 R_{t'}^a + w_3 P_{t'}^a)}} \quad \text{(equation 5)}$$

Although applying the previously mentioned Prospect Theory to the basic QR model did not significantly improve performance, a great deal of theoretical and empirical research still suggests humans do not accurately perceive probabilities. This research provides insight on reasons the original application may not have been successful and also potential ways to improve the SUQR model. Prospect Theory requires the adjustment of both probabilities and values. Given the SUQR model does not directly pair probability with values, we limit our discussion in this research to probability transformations. In Prospect Theory, the probability weighting function must follow an inverse S-shape to account for the Allais paradox (Birnbaum 2006). While this shape has been widely observed in subsequent research, many recent studies have also found S-shaped probability curves at the aggregate level (Etchart-Vincent 2009). Additionally, the later work improving upon Prospect Theory, known as Cumulative Prospect Theory, suspects "decision weights may be sensitive to the formulation of the prospects, as well as to the number, the spacing and the level of outcomes" and the cumulative functional "unlikely to be accurate in detail" (Tversky and Kahneman 1992). In the security games, attackers are forced to pick between a minimum of 8 gates while most previous research experiments just present two or three options. Additionally, attackers face a large combination of outcome levels and spacing with each possible choice. Given the different research observations on

probability transformations, a probability weighting function that allows for both S-shapes and inverse S-shapes seems most appropriate for work in the SSG domain.

The most popular specification of the probability weighting function contains parameters that allow for adjustment of both curvature and elevation (Etchart-Vincent 2009). The specifics of this function originally suggested by Goldstein and Einhorn (1987) will be discussed in the next section.

Model

We expanded the original SUQR model to account for attackers' perceptions of target coverage probabilities by including a probability weighting function. In order to combine the SUQR model and a probability weighting function, we began with the basic SUQR model which includes the previously introduced 3-parameter utility function (equation 4). We replace the x_t representing actual defender coverage probability with $p(x_t)$ representing perceived defender coverage probability, as seen in equation 6.

$$q_t = \frac{e^{\lambda(w_1 p(x_t) + w_2 R_t^d + w_3 P_t^d)}}{\sum_{t'} e^{\lambda(w_1 p(x_{t'}) + w_2 R_{t'}^d + w_3 P_{t'}^d)}} \qquad \text{(equation 6)}$$

Equation 7 illustrates the 2-parameter probability weighting function to calculate $p(x_t)$.

$$p(x_t) = \frac{\delta x_t^{\gamma}}{\delta x_t^{\gamma} + (1-x_t)^{\gamma}} \qquad \text{(equation 7)}$$

The δ term governs elevation and gives information about where the attraction level of the gamble changes from optimism to pessimism or vice versa. The γ term governs curvature and gives information about discriminability between probabilities. As γ approaches 1, the curve becomes flatter, and the individual can more easily discriminate between probabilities. Additionally, $\gamma > 1$ generates the S-shape and $\gamma < 1$ generates an inverse S-shape (Etchart-Vincent 2009).

In order to learn parameters, both for the basic SUQR model and our expanded SUQR model, we employed maximum likelihood estimation (MLE) (Hastie, Tibshirani, and Friedman 2009). We attempt to model the decision making of the general population vice individuals as insufficient data exists for specific individuals. Additionally, we set $\lambda = 1$ without loss of generality.

For the simple model with a set of J games each with a set of T targets and the given defender strategy x, we represent the log-likelihood of (w_1, w_2, w_3) in equation 8. N_t represents the number of subjects attacking target t.

$$log\ L\ (w_1, w_2, w_3 | x) = \sum_{j=1}^{J} \sum_{t=1}^{T} N_t\ log\ [q_t\ (w_1, w_2, w_3 | x)] \qquad \text{(equation 8)}$$

As demonstrated in previous work, this function can be shown as concave and having a unique local maximum point (Nguyen et al 2013). After separating our real-world observations into training, validation and test sets, we used Microsoft Excel's Generalized Reduced Gradient (GRG) nonlinear solver function to obtain the weights(w_1, w_2, w_3) for the basic SUQR model.

For the expanded model with a set of J games each with a set of T targets and the given defender strategy x, we represent the log-likelihood of (γ, δ, w1, w2, w3) in equation 9, where δ and γ are ≥ 0. N_t represents the number of subjects attacking target t.

$$log\ L\ (\gamma, \delta, w_1, w_2, w_3 | x) = \sum_{j=1}^{J} \sum_{t=1}^{T} N_t\ log\ [q_t\ (\gamma, \delta, w_1, w_2, w_3 | x)] \qquad \text{(equation 9)}$$

Unlike the simple model, the expanded model is nonlinear and convex making it extremely difficult to solve for a global maximum. Consequently, we focused on methods to find local optima that have increased probability of being a global solution. First, we utilized Microsoft Excel's GRG nonlinear solver to find an initial local optimum for each of the training sets. Then, the initial solutions of these parameters helped us set bounds to conduct a more robust search for a solution. Given the maximum and minimum of our initial solutions, we selected inclusive bounds for δ and γ of 0 to 150 and inclusive bounds for each of the weights of -50 to 50. Next, we employed the multi-start feature to run the GRG nonlinear solver from 100 different random starting points. This provides an increased level of coverage within the bounds and results in the best selection of several locally optimal solutions. No constraints were found to be binding during this process. While this method does not guarantee a global solution, the probability of reaching it certainly increases.

Model Evaluation Methodology

We performed model validation to be able to better assess how our results will apply to an independent data set. Specifically, we employed 5-fold cross validation to evaluate the parameters we found in each model. This method optimizes the model across all the data while minimizing any bias. The folds for training and validation each contained 7 games. The test set consisted of 9 randomly selected separate games. After learning the parameters from the training sets, we compared them against the validation sets. Finally, we selected the parameters providing the smallest MLE value in validation and then applied them to the testing data to determine the model's goodness of fit on an independent data set.

We then calculated a Mean Absolute Deviation (MAD) and a Mean Squared Error (MSE) in order to compare the two models. This allowed us to quantify how close the predicted response was to the observed real-world response. We compared each of these values for both models assuming the model providing lower values would be better.

In order to directly compare the basic SUQR with our expanded model, we used a likelihood ratio test (LRT) to determine if the difference in value is statistically significant (Neyman and Pearson 1933). This test can be used to compare a simple model nested in an expanded model. First, we establish our null and alternative hypothesis as below.

$$H_0 : \gamma = \delta = 1$$
$$H_A : \gamma, \delta \neq 1$$

Then, we calculate the test statistic. Equation 10 represents the basic formula for the ratio test. The likelihoods of the simple model and expanded models are depicted as $\mathcal{L}_{simple}$ and $\mathcal{L}_{expanded}$ respectively.

$$LRT = -2 \ln\left(\frac{\mathcal{L}_{simple}}{\mathcal{L}_{expanded}}\right), \quad \text{(equation 10)}$$

The quotient rule of the natural logarithm allows the basic equation to be transformed as seen in equation 11.

$$LRT = -2\left(\ln(\mathcal{L}_{simple}) - \ln(\mathcal{L}_{expanded})\right) \quad \text{(equation 11)}$$

With a large sample size, this test statistic approximates χ^2 with the degrees of freedom equal to the difference in the number of parameters between the two models (Wilks 1938). For our model comparisons, the degrees of freedom equal 2. The test statistic can then be used to find a corresponding p-value using a χ^2 look-up table.

Experiment Data and Results

Our data is described in the introduction. We applied both a simple and an expanded variation of the SUQR model to the data in order to compare which provides a better approximation of perceived coverage probabilities when compared to the actual data from AMT workers.

Figure 2 displays parameter estimation results and model metrics following a 5-fold cross validation of the expanded and simple SUQR models for the 8-gate data. In 4 out of the 5 validation rounds, the expanded model resulted in an improved MAD, MSE, and MLE. When evaluating the test data, the simple model produced slightly better MAD and MSE values and the expanded model produced a better MLE value. Next, we conducted the LRT for the 8-gate game data. We calculated a test statistic of 3.61 which corresponds to a p-value of 0.16. This means we must fail to reject the null hypothesis at the 90% and 95% significance level. So for the 8 gate game, we cannot say the expanded model provides significant improvement over the simple model.

Testing

Simple	Expanded
1.00	2.86
1.00	16.65
-14.01	-5.87
0.49	0.47
0.30	0.30
0.0392	0.0427
0.0040	0.0051
-510.38	-508.58

Figure 2.

Figure 3 displays parameter estimation results and model metrics following a 5-fold cross validation of the expanded and simple SUQR models for the 24-gate game data. In all five validation rounds, the expanded model had better values for MAD, MSE, and MLE than the simple model. After evaluating the test data, we found the expanded model also outperformed the simple model in regards to MAD, MSE and MLE. We also conducted an LRT for the 24-gate game data and got a value of 27.06 which corresponds to a p-value of 0. This led us to reject the null hypothesis at the 90% and 95% significance level. For the 24-gate game, we can conclude that the expanded model provides significant improvement over the simple model.

	Simple	Expanded
Testing	1.00	3.91
	1.00	67.32
	-15.94	-5.83
	0.61	0.57
	0.30	0.28
	0.0372	0.0359
	0.0020	0.0019
	-787.07	-773.54

Figure 3.

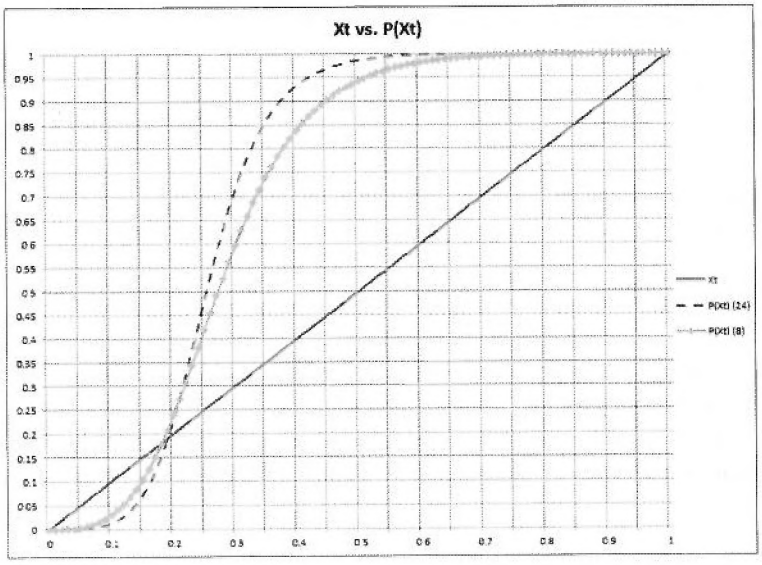

Figure 4.

Figure 4 shows a graphical representation of the perceived probability weighting function (equation 7) using our optimal parameters. The blue line (*Xt*) represents actual coverage probability ranging from 0 to 1. The dotted green line (P(Xt) 8) represents the 8-gate perceived probabilities using γ of 2.86 and δ of 16.65. The dashed red line (P(Xt) 24) represents the 24-gate perceived probabilities using γ of 3.91 and δ of 67.32. The point where these lines change from underweighting coverage probability to overweighting coverage probability is approximately 0.18. Both curves signify that people tend to underestimate low probabilities and overestimate high probabilities.

In the context of security games, the results of the model show small increases in defender coverage probability at certain points can lead to larger decreases in the probability a defender will attack. Specifically, the 24-gate model shows increasing defender coverage 10% at a gate with an average coverage level of 29% will actually reduce the average likelihood of attack by 14% (from 18% to 4%).

Conclusions

Our results from the 24 gate game show the expanded model can provide significant improvement with respect to the simple model. In these cases, defenders should be able to seize upon the increased knowledge of a point where adversaries begin to overestimate defender coverage and improve their strategies through a reallocation of security resources. In other cases, the results from the 8 gate game still suggest the expanded model is unlikely to substantially reduce predictive accuracy. The increased computational requirements of the expanded model should still be considered before deployment in new environments.

At this point, our evidence does not allow us to draw definitive conclusions regarding other cases where the expanded model will result in improvement. Still, the increase in the δ and γ parameters between the 8 and 24 gate games may provide possible insight for future work. Given the rewards and penalties for the games were all taken from the same distribution, those games with increased number of gates may result in more extreme probability distortions and be better suited to the expanded model.

References

Birnbaum, Michael H. (2006) "Evidence against Prospect Theories in Gambles with Positive, Negative, and Mixed Consequences." *Journal of Economic Psychology* 27.6: 737-61. *ScienceDirect.*

Etchart-Vincent, Nathalie. (2009) "Probability Weighting and the 'level' and 'spacing' of Outcomes: An Experimental Study over Losses." *Journal of Risk and Uncertainty* 39.1: 45-63.

Goldstein, William M. and Hillel J. Einhorn. (1987) "Expression Theory and the Preference Reversal Phenomenon." *Psychological Review 94*: 236-254.

Hammond, Kenneth R. (1955) "Probabilistic Functioning and the Clinical Method."*Psychological Review* 62.4: 255-62.

Hastie, T., R. Tibshirani, and J. Friedman. (2009) The elements of statistical learning 2nd edition.

Kahneman, D. and A Tvesky. (1979) "Prospect theory: An analysis of decision under risk." *Econometrica*, 47.2 :263–292.

Luce, R. D. (1959) Individual Choice Behavior: A Theoretical Analysis. New York: Wiley

Mckelvey, Richard D. and Thomas R. Palfrey. (1995) "Quantal Response Equilibria for Normal Form Games." *Games and Economic Behavior* 10.1: 6-38.

McFadden, Daniel. (1976) "Quantal Choice Analysis: A Survey." *Annals of Economic and Social Measurement* 5.4: 363-90.

Neyman, J and E. S. Pearson. (1933) "On the Problem of the Most Efficient Tests of Statistical Hypotheses." *Philosophical Transactions of the Royal Society of London. Series A. Containing Papers of a Mathematical or Physical Character* 231: 289-337

Nguyen, Thanh H., Rong Yang, Amos Azaria, Sarit Kraus, and Milnd Tambe. (2013) "Analyzing the Effectiveness of Adversary Modeling in Security Games." Association for Advancement of Artificial Intelligence (AAAI) Conference.

Pita, J., M. Jain, F. Ordonez, C. Portway, M. Tambe, C. Western, P. Paruchuri, and S. Kraus. (2008) "Deployed armor protection: The application of a game theoretic model for security at the los angeles international airport." International Joint Conference on Autonomous Agents and Multi-Agent Systems.

Tversky, Amos, and Daniel Kahneman. (1992) "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty* 5: 297-323.

Wilks, S. (1938) "The Large-Sample Distribution of the LIkelihood Ratio for Testing Composite Hypotheses." *The Annals of Mathematical Statistics*. 9.1: 60-62.

Yang, R.; Kiekintveld, C.; Ordonez, F.; Tambe, M.; and John, R. (2011) "Improving resource allocation strategy against human adversaries in security games." *IJCAI*: 458–464.